

2

DTIC FILE COPY

*A Workshop on
"High-Level Connectionist Models"*

FINAL REPORT

April 9-11, 1988
Las Cruces, New Mexico

Jordan Pollack
John Barnden
Computing Research Laboratory
New Mexico State University
Las Cruces, NM 88003

AD-A198 759

Table of Contents

DTIC
ELECTE
AUG 08 1988
S D
E

1. PROPOSAL SUMMARY	1
1.1. The Problem	1
1.2. The Participants	2
1.3. The Organization	2
1.4. The Setting	3
2. PROCEEDINGS	4
2.1. Introduction	4
2.2. Presentation Summaries	4
2.3. Discussion & Common Themes	7
2.4. Conclusion	8
3. PARTICIPANT LIST	9
4. SCHEDULE OF EVENTS	10
5. ABSTRACTS OF PRESENTATIONS	11
6. REFERENCES	17

This work relates to Department of Navy Grant N00014-88-J-1066, issued by the Office of Naval Research.

THIS DOCUMENT HAS BEEN APPROVED
FOR PUBLIC RELEASE AND IS NOT
WITHSTANDING IT IS UNCLASSIFIED, UNCONTROLLED

88 6 3 007

A Workshop on "High-Level Connectionist Models" FINAL REPORT

April 9-11, 1988
Las Cruces, New Mexico

Jordan Pollack
John Barnden
Computing Research Laboratory
New Mexico State University
Las Cruces, NM 88003



Accession For	
NTIS GRANT	A
DTIC TAB	□
Unannounced	□
Justification	A
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

1. PROPOSAL SUMMARY

1.1. The Problem

Connectionist models have been very successful in low level cognitive modeling of perception, motor control, associative memory and even some elements of language processing. As attempts are made to build models of high-level cognitive processes that have traditionally been handled with sequential symbolic computing — such as reasoning, planning and natural language understanding — certain technical difficulties concerning processing and representation continually arise. A few researchers have been attempting to overcome these difficulties without resorting to massive duplication of resources or external serial logical control. Despite their efforts, however, only limited progress has been made toward solving the problems.

Yet certain elements of connectionism (e.g. associative memories, learning, distributed representations) seem to provide qualities of robustness, flexibility and graceful degradation that are difficult to achieve in the conventional, symbolic framework. We are therefore faced with the problem of determining how connectionism can fulfill this promise while achieving the power of traditional symbolic processing in a reasonably natural way.

→ The difficulties that continually show up in connectionist modelling attempts directed towards high-level cognitive processing and knowledge representation include the inter-related problems of generative capacity, representational adequacy, variable binding, multiple instantiation of schemata (concepts, frames, etc.), rapid construction and modification of information structures, task control, and recursive processing. These problems are of more direct concern to computer scientists than to psychologists, philosophers or neuroscientists. Accordingly, the purpose of the workshop is to bring together computer-science oriented connectionist researchers who have addressed the problems, so as to understand the issue of achieving high-level cognitive processing in connectionist systems more clearly, and to pass the results on to the scientific community at large.

→ The following is a list of preferred topics for presentations and discussions at the workshop. The topics are not orthogonal.

1. connectionist/neural implementation of (aspects of) commonsense reasoning, planning, natural language understanding and rapid, "one-shot" learning, → neural

- 2 ways of coping with the productivity of natural language
- 3 ways of encoding complex information in connectionist/neural models
- 4 capabilities and limitations of coarse-coded and distributed representations
- 5 the relative advantages of value-unit, variable-unit and intermediate representations
- 6 the relative advantages of localist and distributed representations
- 7 connectionist implementation of temporary, dynamic, complex data structures (such as frame instantiations, network fragments, stacks, trees)
- 8 multiple simultaneous instantiation of rules, concept structures, schemata, frames, etc.
- 9 variable binding
- 10 task control, including sequencing, iteration and recursion
- 11 real and apparent rule following. (1.18) ←

1.2. The Participants

The list below indicates the group of scientists who will be attending the workshop. A few observers will also be present, both local faculty or students as well as representatives of the supporting agencies.

Host	Larry Birnbaum	Yale University
	Larry Bookman	Brandeis University
	John Barnden	New Mexico State University
	Garrison Cottrell	University of California, San Diego
	Mark Derhick	Carnegie Mellon University
	Joachim Diederich	ICSI, University of California, Berkeley
	Michael Dyer	University of California, Los Angeles
	Jeffrey Elman	University of California, San Diego
	James Hendler	University of Maryland
	Wendy Lehnert	University of Massachusetts
Host	Jordan Pollack	New Mexico State University
	Roger Schvaneveldt	New Mexico State University
	Lokendra Shastri	University of Pennsylvania
	David Touretzky	Carnegie Mellon University

1.3. The Organization

The meeting will be kept as informal as possible within the following constraints.

- The meeting will consist mainly of presentations by the participants followed by open critical discussions.
- Participants will be encouraged to present work that is not already familiar to the others.
- Each presentation together with following discussion will fit in a time slot of approximately 60 minutes. The presentations themselves will be confined to 40 minutes each.
- Time periods will also be reserved for further discussion on topics suggested during the workshop by the participants.
- Each of these discussions will be moderated by one of the organizers or by a participant appointed during the workshop. The moderator will be responsible for keeping the discussion on track and for suggesting points for debate when appropriate.

- The workshop will be audiotaped and later transcribed. The transcriptions of discussions will be edited into the workshop proceedings.

If, therefore, there are 5 presentations on each of the three days, five hours a day will be taken up by presentations and the following discussions. This will leave time for a 40-minute lunch break, two other 30-minute breaks, and an hour of extra discussion. Altogether, each day's preplanned events will occupy eight and a half hours (say from 9:00am to 5:30pm).

When Barnden joined CRL in August 1987, he had already contracted with ABLEX to edit a book series entitled *Advances in Connectionist and Neural Computation Theory*, whose topic area has a large overlap with that of the workshop. Therefore, the plan of using the first volume in the series as a destination for the workshop papers was conceived. The proceedings will contain papers by the participants along with edited transcripts of workshop discussions — both the debates generated by specific presentations and the more general discussions.

1.4. The Setting

The workshop will take place Saturday through Monday, April 9-11, 1988 in the Holiday Inn de Las Cruces, with accommodations in a neighboring suite-only hotel. Las Cruces, the largest city in Southern New Mexico, and home to New Mexico State University, is situated in a valley between the spectacular Organ Mountains and the Rio Grande River. It is located one hour from El Paso, Texas, and from the White Sands National Monument, and two hours from the Ruidoso and Cloudcroft skiing areas. Weather in April is in the 60's or 70's during the day, dropping to the 40's by night.

Previous workshops and conferences held by the Computing Research Laboratory include a workshop on Pyramidal Computers hosted by Leonard Uhr in March 1985, the Foundations of AI workshop hosted by Derek Partridge in February 1986, a meeting on Graph Theory and Computer Science hosted by Frank Harary in April 1986, and TINLAP3 hosted by Yorick Wilks in January 1987.

The Computing Research Laboratory was founded in July 1983 as an autonomous unit in the College of Arts & Sciences of New Mexico State University. With initial funding provided by the state legislature, CRL has developed into a first-class center for research in artificial intelligence and cognitive science. Specific areas include: The human-computer interface, natural language understanding, knowledge representation and reasoning, connectionism, robotics, and computer vision. The Laboratory currently employs, amongst others, a full-time director, fourteen faculty members with joint appointments in the departments of Computer Science, Electrical Engineering, Mathematics and Psychology, eight full-time researchers and over thirty research assistants.

2. PROCEEDINGS

2.1. Introduction

The intention of this workshop was to bring together a small group of computer scientists to focus on the interaction between AI and connectionism. The two fields are often posed as paradigmatic enemies (See (Dreyfus & Dreyfus, 1988), and other papers in the recent issue of *Daedalus*), and there seems to be a risk of severing them. Few connectionist results are published in the mainstream AI publications and conferences other than *Cognitive Science*, and many of neural-network researchers and industrialists proceed without due consideration of the problems (and progress) of AI.

The workshop was assembled by invitation, and we strived both for geographic and ideological breadth. The questions posed to the participants involved the application of connectionism to traditionally symbolic tasks, approaches to overcoming some of the limitations exhibited by connectionist models of the past several years, and opportunities for connectionism to make contributions to the study of cognition.

The participants were Wendy Lehnert, Michael Dyer, and James Hendler (as AI folk who are testing the connectionist waters), Garrison Cottrell and Jeffrey Elman (representing the hard-line PDP school), Lokendra Shastri and Joachim Diederich (the Feldman School), Mark Derthick and David Touretzky (the Hinton School), Lawrence Bookman (the Waltz School) and Lawrence Birnbaum (the Loyal Opposition). In addition, Jordan Pollack, John Barnden and Roger Schvaneveldt represented their own school, New Mexico State University.

Each participant was given exactly an hour to present their views and manage the oftentimes vociferous audience. The 14 presentations were squeezed into 2 & 1/2 days, accompanied by long, social evenings. The workshop was followed by a visit to White Sands National Monument.

What follows below is a summary of each person's presentation (in alphabetical order), then a discussion of some of the major themes emerging from the workshop.

2.2. Presentation Summaries

John Barnden, of New Mexico State University, described the current status and future plans for his system, called CONPOSIT (Barnden, 1986, 1987). Implemented on Nasa's MPP, the system makes use of a two-dimensional array of registers constructed of simple neuron-like processors and represents complex symbolic structures, using adjacency between registers as a form of variable-binding. He demonstrated how production rules are currently implemented and timed, and proposed how a large set of such rules could be managed. Barnden also showed that his system could reason syllogistically, directly implementing Johnson-Laird's notion of a mental model (Johnson-Laird & Bara, 1984).

Lawrence Birnbaum, of Yale University, spoke about hard problems in planning and understanding which would not just go away for connectionists. His two examples involved long-chained inferences over very abstract schemata. One was a story of a detective following a suspect into a hardware store. After being spotted by the suspect, he buys a bucket. The second was a hypothetical (at the time) story about Dukakis worried about winning the NY primary without alienating the Black Vote. Gore attacks Jackson's views on the Mideast, and Dukakis carries the state. Both "looking inconspicuous" and "helping" (by acting as a shill) involve reasoning about an agent's own internal state, and thus very strong self-referential representational systems, at which connectionism flounders.

Lawrence Bookman, of Brandeis University, argued that connectionism needed better forms of modularity in order to scale up. He described his work on Network

Regions (Chun et. al., 1987), which are a hierarchal abstraction for spreading activation networks. A network region is an abstract unit standing for a collection of units, whose input, activation level, and output are averages over the corresponding values for those units. He showed how interacting clusters of winner-take-all networks used for parsing and lexical disambiguation, such as those used by (Waltz & Pollack, 1985), could be more cleanly implemented in his system. A long discussion ensued about such super-nodes and the actual costs of their implementation.

Garrison Cottrell, of the University of California at San Diego, described his recent work on learning to compress gray-scale images (Cottrell et. al., 1987). That system learned a linear transformation of a well-known algorithm called Principal Components analysis. He then described some extensions of this work towards using the internal representations as a basis for semantics for words and on various projects involving the use and representation of sequences of images. images.

Mark Derthick, of Carnegie-Mellon University, described the genesis of his system called μ KLONE (Micro-KL-one), which is an energy-minimization approach to generating a finite model of a first-order logic theory. He argued that his approach was better than a theorem prover because it degrades more gracefully when resource limitations prevent running to completion. His system also provides a formal probabilistic theory for common-sense reasoning in the face of a knowledge base that may be logically incomplete or inconsistent. He described the simplifications and modifications necessary to avoid intractability in his system.

Joachim Diederich, of the new International Computer Science Institute at Berkeley, described his research on spreading activation and various forms of network-propagation (Diederich 1986, 1987, 1988). He contrasted "pulse-specific" discrete marker-passing with "source-specific" analog spreading activation in several network taxonomies, and the problems for each. Using psychological evidence, he argued for a form of inhibition (Renshaw) over the usual winner-take-all lateral inhibition both as less resource intensive (requiring $O(n)$ rather than $O(n^2)$ links) and as a better account of classification data from humans under stress. Renshaw-inhibition is one of the most important inhibition patterns in animal nervous systems. The presented application of Renshaw-inhibition avoids false classification by use of intermediate units between concept units and property units in a spreading activation network. If two or more concepts share a property, and both concepts are source units for the spreading activation process, only the strongest concept unit and the particular property unit remain active simultaneously. Diederich also described the organization and work-in-progress of ICSI, which is headed by Jerry Feldman.

Michael Dyer, of the University of California at Los Angeles, laid out his general philosophy on "Symbolic Neuro-Engineering", through which traditional AI tasks are decomposed into modules which can be implemented using connectionist techniques, yielding more robust and flexible systems. He described a number of research projects he and his colleagues were engaged in. He demonstrated how an extended back-propagation scheme discovered representations of lexical entities while the network was performing the (McClelland & Kawamoto, 1986) case-frame mapping task (Miikkulainen & Dyer, 1988). As a result, a microfeature specification for the input is no longer necessary. Dyer also described DUAL (Dyer et. al., 1988), a PDP architecture able to represent labelled, directed graphs with cycles. The general technique manipulates the entire weight matrix formed by one network as a pattern of activation in a larger network. As a semantic network is encoded in the PDP architecture, distributed representations are formed for each node in the semantic network, with structurally similar nodes developing similar patterns of activation. Dyer only briefly hinted at work on: (a) associating visual motion with linguistic descriptions, which he referred to as the grounding problem (Nenov & Dyer, 1988), (b) variable binding techniques using conjunctive coding (Dolan & Dyer, 1987, 1988), (c) localist architectures for

dynamic word sense reinterpretation (Sumida et. al., 1988) and (d) a model of completely parallel language generation (Gasser & Dyer, 1988).

Jeffrey Elman, of the University of California at San Diego, presented his latest work on learning to process and predict sequences (Elman, 1988). He uses a recurrent form of back-propagation (Rumelhart et. al., 1986a), where the output of units in a hidden layer are directly fed back into input units. Using very long input sequences, Elman trains a network in prediction tasks. For example, one task was to predict the next word in a long concatenated sequence of short, grammatically-generated sentences. In learning to do so, the network devised a classification for the words (into noun/verb, animate/inanimate, etc.) which was revealed through cluster analysis. Another network learned to resolve pronouns in a large body of grammatically-generated variable-length sentences.

James Hendler, of the University of Maryland, discussed his work on a hybrid model of marker-passing which descends into the subsymbolic, using microfeatures to mediate similarity. The problem was to recognize, in a planning domain, that a letter-opener is similar enough to a knife to set off alarm at an airport security gate. He showed how this could be accomplished, without a direct link to WEAPON, by spreading through microfeatures such as METAL, POINTY, etc. He proposed the next generation of his system, which will develop the microfeatures automatically using back-propagation. Hendler also took on the whole philosophical question of AI versus Connectionism. Two interesting analogies were discussed. Geoff Hinton was quoted as comparing AI to zoology (descriptive) and Connectionism to molecular biology (constructive). Hendler pointed out the appropriateness of this mapping but cautioned that cognition is much more like evolution than autopsy. John McCarthy was quoted as saying that AI and Connectionism were two horses in the same race, and they should be permitted to run the course; Hendler asserted that there is really only one horse, but it has two ends.

Wendy Lehnert, of the University of Massachusetts, made a strong case for hybrid modeling: taking the best of both worlds for building useful systems. She described her new sentence analysis system, which uses a stack and a copy mechanism for control, marker-passing for prediction, and numerical spreading-activation for smooth decision-making. She showed how the system could give a unified syntactic and semantic account for parsing and prepositional phrase attachment and flexibly implement a "no-crossing of branches" constraint.

Jordan Pollack, of New Mexico State University, argued that multiplicative connections would lead to more powerful models, that stability and complexity problems could be handled by programming conventions, and described two such conventions. He also described recent work on devising compositional distributed representations for variable-sized stacks and trees using a recursive form of auto-association (Pollack, 1988). He argued that, in addition to overcoming major complaints about the adequacy of connectionist representations such as those leveled by (Fodor & Pylyshyn, 1988), the representations developed by his system could be the basis of systems which perform inference using fixed-width pattern association.

Roger Schvaneveldt, also of New Mexico State University, discussed his attempt at replicating the schema system of (Rumelhart et. al., 1986b), using a matrix of content-features which represent prototypical rooms. He compared and contrasted the PDP dense relaxation system with hierarchical clustering techniques, and with his PATH-FINDER graph-theoretic algorithm (Schvaneveldt et. al., In Press)

Lokendra Shastri, of the University of Pennsylvania, took on the question of what contribution connectionism could make to AI. He described his work on knowledge representation (Shastri, 1988) and reasoning in a more general framework of discovering fast and tractable means of performing inference. Shastri stressed that robust intelligent behavior requires an extremely fine-grained decomposition of knowledge, and that

connectionism forces us to discover such decompositions. One success has been the complete compilation of default reasoning using property inheritance and classification into a spreading activation connectionist network.

David Touretzky, of Carnegie-Mellon University, raised a warning flag about the abundance of trivial models coming out of connectionism. He said that certain ideas had already reached their limit of usefulness and could stagnate the field. Continued attempts to reduce cognition to lateral-inhibition among localist units, rote associative memory with no inferencing ability, and symbol transformation by 3-layer back-propagation nets are all in this class. He presented several projects that he has been working on, including the DUCS system for representing frames (Touretzky & Geva, 1987), a system for prepositional phrase attachment, and a new ambitious project on connectionist metonymic reasoning.

2.3. Discussion & Common Themes

Each participant presented their results and views while being constantly interrupted with criticisms and tangential discussions. Some recurrent themes did, however, emerge from this chaos.

Pokey connectionist Models

Originally Birnbaum's colorful complaint, the idea of new connectionist implementations of things which LISP can do easily, such as CONSing, Towers of Hanoi, syntactic parsers, production systems, logical operations, etc., generally met with negative reactions from the audience. Such implementations are only interesting when they demonstrate useful capabilities beyond simple symbolic adequacy.

Hybrid Modeling

Both Wendy Lehnert and James Hendler argued strongly in favor of mixing symbolic and connectionist methodologies. Although such work proceeds and is useful, questions about the utility and epistemology of such models remain. There was a general lack of mutual understanding on the notion of discrete abstract "levels" in cognitive models, or on the notion of compilation between levels. A good question to ask of a hybrid model is if it truly bridges levels from pre- to post-symbolic processing. Dyer argued that work should proceed on at least four levels, including the knowledge level, activation-spreading semantic networks, connectionist, and neural, with an eye towards the key questions of how each level is justified by the one above and can be embedded into the one below.

Moving-Target Learning

In most systems using back-propagation, the training environment consists of a stable set of input and target patterns. An interesting commonality emerged in the work of Dyer, Elman, and Pollack: In each of these systems, using various recurrent forms of back-propagation, the learning environment changes along with the weights in the networks. Given that semantic networks, word sequences, and trees, respectively, are among the first dynamically sized complex data-structures to be learned, represented and processed using connectionist networks, the moving-target strategy may be a win.

Connectionism may redefine *Symbol*

In a very lively discussion period, Schvaneveldt brought up the whole question of the meaning and use of symbols, and Elman complained that they were just useless baggage. In AI, symbols have no "internal structure" and thus mean very little, they are just used as names or pointers to larger structures of symbols, which are reasoned with (slowly). The subsymbolic distinction was considered incoherent (and Smolensky was not present to defend it), because just about every connectionist model discussed was symbolic in nature, even at the so-called "micro-feature" level. The essential difference between the early neural network research

and modern connectionism is that AI has happened in-between them. There was some suggestion that a new kind of symbol ("symboid", according to Dyer) might emerge from connectionism. For example, a reduced representation could be considered a such a symboid, given that it can "point" to a larger structure through a reconstruction algorithm. Symboids may have an advantage over symbols, in that they possess internal structure which can be reasoned about.

Hard Problems Remain Hard

A definite consensus was that, while a change in computational methodology might make some problems easier for connectionist or symbolic processing, the hard problems are not going to go away, and thus intelligence will not be solved either by three-layer back-propagation, or by physicists studying dynamical systems. The list of hard problems discussed include the organization and encoding of knowledge, inference, managing dynamic instantiation of structures, and rapid (one-shot) learning. Natural language, planning, and reasoning all require infinite generative capacity and the ability to combine old knowledge in new ways. In order to approach these problems, connectionists need to continue work on complex representations and on the effective use of modularity and hierarchy.

2.4. Conclusion

Given the diverse backgrounds and methodologies of the participants, and the brief and intense nature of the workshop, it cannot be said that true consensus was reached or new scientific ground was broken. Each participant will undoubtedly continue on their own path, influenced, perhaps, by the common experience.

However, it must be considered a great success from the point of view of communication. We placed these people in close quarters for an extended period of time, and, despite a lot of heated argument, criticism, slicing, and dicing, nobody was induced to suicide. It is safe to say that research in connectionist approaches to higher-level cognitive functions will continue for some time.

3. PARTICIPANT LIST

John Barnden
Box 3CU
New Mexico State University
Las Cruces, NM 88003
jbarnden@nmsu.edu
(505) 646-6235

Lawrence Birnbaum
Department of Computer Science
PO Box 2158 Yale Station
New Haven, CT 06520
birnbaum@yale.edu
(203) 432-1223

Lawrence Bookman
Computer Science Department
Brandeis University
Waltham, MA 02154
(617) 736-2704

Garrison Cottrell
Computer Science & Engineering C-015
University of California
La Jolla, CA 92093
gary%cs@ucsd.edu
(619) 534-6640

Joachim Diederich
Int'l Computer Science Institute
1947 Center Street
Berkeley, CA 94704
joachim@icsi.berkeley.edu
(415) 643-9153

Mark Derthick
Computer Science Department
Carnegie Mellon University
Schenley Park
Pittsburgh, PA 15213
derthick@cs.cmu.edu
(412) 268-3066

Michael Dyer
Computer Science Department
3532 Boelter Hall
University of California
Los Angeles, CA 90024
dyer@cs.ucla.edu
(213) 206-6674

Jeffrey Elman
Department of Linguistics, C-008
University of California
La Jolla, CA 92093
elman@amos.ling.ucsd.edu
(619) 534-1147

James Hendler
Department of Computer Science
University of Maryland
College Park, MD 20742
hendler@mimsy.umd.edu
(301) 454-4148

Wendy Lehnert
Computer and Info. Science Dept.
University of Massachusetts
Amherst, MA 01003
lehnert.umass@csnet.relay
(413) 545-3639

Jordan Pollack
Box 3CRL
New Mexico State University
Las Cruces, NM 88003
(505) 646-5861
pollack@nmsu.edu

Roger Schvaneveldt
Box 3452
New Mexico State University
Las Cruces, NM 88003
(505) 646-1047
schvan@nmsu.edu

Lokendra Shastri
Dept. of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104
shastri%cis.upenn.edu
(215) 898-2661

David Touretzky
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
touretzky@cs.cmu.edu
(412) 268-7561

4. SCHEDULE OF EVENTS

Time	Saturday	Sunday	Monday
8:00-8:30	Introductory Remarks		
8:30-9:30	Diederich	Hendler	Touretzky
9:30-10:30	Bookman	Lehnert	Pollack
10:30-11:00	COFFEE BREAK	COFFEE BREAK	COFFEE BREAK
11:00-12:00	Elman	Cottrell	DISCUSSION FAREWELL
12:00-1:00	LUNCH Holiday Inn	Barnden	LUNCH Holiday Inn
1:00-2:00	Schvaneveldt	BRUNCH	
2:00-3:00	Dyer	Double	
3:00-3:30	BREAK	Eagle	
3:30-4:30	Birnbaum	Derthick	
4:30-5:30	DISCUSSION	Shastri	
5:30-8:00	FREE DRINKS Plaza Suites	RECEPTION Roger & Ann Schvaneveldt's	
8:00-11:00	DINNER Hacienda	DINNER Tatsu	

5. ABSTRACTS OF PRESENTATIONS

Saturday, 8:30-9:30

*Knowledge Representation and Recruitment Learning
in a Connectionist Inheritance Network.*

Joachim Diederich
ICSI, Berkeley

A connectionist knowledge representation system is introduced to realize reasoning corresponding to multiple inheritance in standard semantic network formalism dealing with evidential information. The system is designed to allow machine learning mechanisms such as classification and recruitment. The knowledge representation has four virtual parts: a concept space for the representation of taxonomic relations between concepts including multiple inheritance paths, an attribute space for the representation of attributes and their values, an instance space for the representation of examples during learning and a free space of uncommitted units which become committed during connectionist recruitment learning.

Saturday, 9:30-10:30

*A Model of Higher-Order Concept Interaction
Using Network Regions.*

Larry Bookman
Brandeis

This talk presents a model of "higher-order" concept interaction through the use of meta-network structures called network regions. Network regions are computational entities that represent the collective state of a set of nodes and hence a "higher-order" abstraction of these nodes. Their interaction with other network components forms our model of concept interaction. Network regions provide more structure for connectionist networks, which can make these networks more modular and thus more understandable; and easier to construct since these components are larger conceptual "chunks." Currently there is no well-defined methodology to define structures within connectionist networks. This severely limits the ability of networks to "scale up" to larger problems. Although current multiple-layered networks provide some form of structure, nodes in these networks do not represent true hierarchical concepts. Regions provide a means of creating hierarchies, as well as shared structures, since they can be embedded and overlapped. In our model, concepts at a particular level need not have explicit knowledge of other concept's subcomponents in order to influence them. Thus network regions can act as building blocks that simplify the construction of large connectionist networks.

Saturday, 11-12:00

Finding Structure in Time.

Jeff Elman
UC San Diego

Time is clearly important in cognition. It is inextricably bound up with many behaviors, such as language, which express themselves as temporal sequences. In serial processing frameworks, the question of how to represent time does not arise. In PDP models, time has historically been converted into space, which leads to many problems. I will discuss an alternative approach which implements the insight that time can be represented by the effect it has on processing and not as an additional input dimension. The approach is rather simple, but the results are sometimes complex and unexpected. Some of the simulations suggest that the problem of time may interact with other problems for connectionist architectures, including the problem of symbolic representation.

Saturday, 1:00-2:00

Proximities, Networks, and Schemata

Roger Schvaneveldt
New Mexico State University

This paper examines the representation of schemata in networks (weighted graphs) and the use of activation to instantiate the schemata in particular contexts. The following questions are of interest: (1) Can direct judgments of co-occurrence provide the basis for such networks? (2) Are the (nearly) complete networks of the connectionist variety essential, or will sparse networks such as Pathfinder networks suffice? (3) How critical are the particular concepts included and the particular activation procedures?

Saturday, 2:00-3:00

Symbolic Neuroengineering

Mike Dyer
UC Los Angeles

The talk will be about (1) levels of analysis for NLP, (2) encoding semantic networks in PDP networks, and (3) eliminating microfeatures.

Saturday, 3:30-4:30

*Complex features in planning and understanding:
Problems and opportunities for connectionism.*

Larry Birnbaum
Yale

Connectionist models have shown that, given certain constraints, perception and interpretation of situations along many different dimensions can be carried out quickly and in parallel. Foremost among these constraints is the limitation to propositional features, i.e., boolean combinations of inputs, at least in models that have been proposed to date. Unfortunately, evidence from natural language processing argues that this restriction is untenable in the long run. Language understanding seems to require explanatory inference processes that are both very quick, and capable of employing highly abstract explanatory structures (e.g., thematic structures), especially in novel cases.

For example, recognizing a novel instance of "helping" entails recognizing the following pattern: "There is a causal chain linking an action of the helper's with a problematic precondition for some goal of the recipient's." This is hard for several reasons. First, the constituents of such concepts are very generic -- almost all of the specificity lies in the structural relations among the elements, rather than the elements themselves. This makes indexing extremely difficult. Second, what was formerly part of our theoretical vocabulary in describing the understanding process ("causal chain") is now attributed to the understander's conceptual repertoire -- that is, understanding such abstract concepts entails reasoning about your own representations. The same problems seem to arise in planning, particularly opportunistic planning, in which an agent adopts goals on the basis of opportunities afforded by the situation in which it finds itself.

However, connectionism may yet triumph. If perception and interpretation really do require complex inference, then there remains the crucial problem of directing and focussing the (comparatively) limited computational resources available. There is a niche, therefore, for theories of attention focussing which do not themselves require complex inference to be carried out. Connectionism seems likely to play an important role here.

Sunday, 8:30-9:30

Intermediate Mechanisms For Activation Spreading

Jim Hendler
University of Maryland

Spreading activation, in the form of computer models and cognitive theories, has recently been undergoing a resurgence of interest in the cognitive science and AI communities. Two competing schools of thought have been forming. One technique concentrates on the spreading of symbolic information through an associative knowledge representation. The other technique has focused on the passage of numeric information through a network. In this talk we show that these two techniques can be merged.

We show how an "intermediate level" mechanism, that of symbolic marker-passing, can be used to provide a limited form of interaction between traditional associative networks and subsymbolic networks. We describe the marker-passing technique, show how a notion of microfeatures can be used to allow similarity based reasoning, and demonstrate that a back-propagation learning algorithm can build the necessary set of microfeatures from a well-defined training set. We discuss several problems in natural language and planning research and show how the hybrid system can take advantage of inferences that neither a purely symbolic nor a purely connectionist system can make at present. Finally, we discuss what this has to say to the notion of connectionism and symbolic AI being separate paradigms.

Sunday, 9:30-10:30

A Connectionist/Symbolic system for Understanding Sentences

Wendy Lehnert
University of Massachusetts

I will describe a natural language sentence analyzer that integrates three distinct architectures: (1) A stack for syntax, (2) marker passing for "predictive semantics," and (3) numerical relaxation for "attachment semantics". I will claim that sentence analysis is best served by heterogeneous architectures rather than homogeneous mechanisms.

Sunday, 11:00-12:00

Procedural Memory in Neural Networks

Gary Cottrell
UC San Diego

Applications of previously developed architectures are proposed for the development of networks that can store and remember sequences of images. The term "procedural memory" is used here to denote the use of the network as a procedure for remembering images, as opposed to a passive storage medium. The application of this to a model of the acquisition of word meanings, both object and event oriented, is discussed.

Sunday, 12:00-1:00

Relative-Position Encoding for Connectionist Data Structures

John Barnden
New Mexico State University

A promising approach to high-level cognitive processing is to encode short-term symbolic data structures in connectionist systems by means of two unusual techniques, called Relative-Position Encoding and Pattern-Similarity Association. These techniques enable a quasi-connectionist system named Conposit to accomplish rule following, variable binding, and construction and analysis of complex, temporary data structures. Conposit is quasi-connectionist in being a computational architecture whose primitive components can readily be implemented connectionistically, although this implementation is merely hinted at here. Conposit is also viewed as a partial and preliminary model of high-level information processing in the brain. Some simulation results are mentioned, in particular for a Conposit version embodying core aspects of Johnson-Laird's "mental model" theory of syllogistic reasoning.

Sunday, 3:30-4:30

Micro-KLONE

Mark Derthick
Carnegie Mellon

I will talk about my connectionist knowledge representation system, micro-KLONE, and how it approximates the ideal behavior specified for rule-based systems. By searching over the space of models, rather than the space of proofs, its actual performance may surpass that of traditional systems when resource limitations are considered.

Sunday, 4:30-5:30

Connectionism and knowledge representation.

Lokendra Shastri
University of Pennsylvania

What do we expect to gain by studying the problem of Knowledge representations and Reasoning within a connectionist framework. Also, what are some important problems that will have to be solved along the way.

Monday, 8:30-9:30

Beyond Associative Memory

Dave Touretzky
Carnegie Mellon

Conventional associative memories aren't powerful enough to account for interesting linguistic behavior because they require an intractable number of training examples. I will describe a particular problem (metonymy) that makes the shortcomings of current connectionist approaches quite clear. Finally, I'll discuss some rough ideas for new specialized architectures for linguistic applications.

Monday, 9:30-10:30

The Solution to Everything

Jordan Pollack
New Mexico State University

The reason I invited all of you here is because, after running into the problems of Generative Capacity, Representational Adequacy, Scaling, and Task Control, and discovering that everybody else crashed into them as well, I figured I'd get some help.

However, I just solved the whole thing, and will talk about Recursive and Distributed, Implausively Compositional, Analog Laminar (RADICAL) Representations, which will allow AI to be performed by fixed-width pattern association and recognition techniques.

You can all go home now.

6. REFERENCES

- Barnden, J. A. (1986). Complex cognitive information-processing: A computational architecture with a connectionist implementation. Technical Report 211, Bloomington: Computer Science Dept., Indiana University.
- Barnden, J. A. (1987). Simulation of an array-based neural net model. In *Proceedings of the First Symposium on the Frontiers of Massively Parallel Scientific Computation*.
- Chun, H. W., Bookman, L. A. & Afshartous, N. (1987). Network Regions: Alternatives to the Winner-Take-All Structure. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*. Milan, 380-387.
- Cottrell, G., Munro, P. & Zipser, D. (1987). Learning Internal Representations from Gray-Scales Images: An Example of Extensional Programming.. In *Proceedings of the Seventh Annual Conference of the Cognitive Science Society*. Seattle, 461-473.
- Diederich, J. (1986). Kognitive Parallelverarbeitung. GMD Report 159, Munich Vienna: R. Oldenbourg Verlag.
- Diederich, J. (1987). *Simulation schizophrener Sprache*. Opladen: Westdeutscher Verlag.
- Diederich, J. (1988). Spreading Activation and Renshaw Inhibition in Neural Networks. In *Proceedings of the Institute of Electrical and Electronics Engineers Second Annual International Conference on Neural Networks*. San Diego.
- Dolan, C. & Dyer, M. G. (1987). Symbolic Schemata, Role Binding and the Evolution of Structure in Connectionist Memories. In *Institute of Electrical and Electronics Engineers First International Conference on Neural Networks*. San Diego.
- Dolan, C. & Dyer, M. G. (1988). Parallel Retrieval and Application of Conceptual Knowledge. UCLA-Artificial Intelligence-88-3, Los Angeles: Artificial Intelligence Laboratory, UCLA.
- Dreyfus, H. L. & Dreyfus, S. E. (1988). Making a Mind versus Modeling the Brain: Artificial Intelligence Again at the Crossroads. *Daedalus*, 117.
- Dyer, M. G., Flowers, M. & Wang, Y. A. (1988). Weight Matrix = Pattern of Activation: Encoding Semantic Networks as Distributed Representations in DUAL, a PDP architecture. UCLA-Artificial Intelligence-88-5, Los Angeles: Artificial Intelligence Laboratory, UCLA.
- Elman, J. L. (1988). Finding Structure in Time. Report 8801, San Diego: Center for Research in Language, UCSD.
- Fodor, J. & Pylyshyn, A. (1988). Connectionism and Cognitive Architecture: A Critical Analysis. *Cognition*, 28, 3-71.
- Gasser, M. & Dyer, M.G. (1988). Sequencing in a Connectionist Model of Language Processing. In *Proceedings of the 12th International Conference on Computational Linguistics*. Budapest.
- Johnson-Laird, P. N. & Bara, B. G. (1984). Syllogistic Inference. *Cognition*, 16, 1-61.
- McClelland, J. & Kawamoto, A. (1986). Mechanisms of Sentence Processing: Assigning Roles to Constituents. In J. L. McClelland, D. E. Rumelhart & the PDP research Group, (Eds.), *Parallel Distributed Processing: Experiments in the Microstructure of Cognition*, Vol. 2. Cambridge: MIT Press.
- Miikkulainen, R. & Dyer, D. G. (1988). Forming Global Representations with Back Propagation. UCLA-Artificial Intelligence-88-7, Los Angeles: Artificial Intelligence Laboratory, UCLA.
- Nenov, V. I. & Dyer, M.G. (1988). DETE: Connectionist/Symbolic Model of Visual and Verbal Association. In *Proceedings of the Institute of Electrical and Electronics Engineers Second Annual International Conference on Neural Networks*. San Diego.
- Pollack, J. B. (1988). Recursive Auto-Associative Memory: Devising Compositional Distributed Representations. MCCS-88-124, Las Cruces: Computing Research Laboratory, New Mexico State University.

- Rumelhart, D. E., Hinton, G. & Williams, R. (1986). Learning Internal Representations through Error Propagation. In D. E. Rumelhart, J. L. McClelland & the PDP research Group, (Eds.), *Parallel Distributed Processing: Experiments in the Microstructure of Cognition*, Vol. 1. Cambridge: MIT Press.
- Rumelhart, D. E., Smolensky, P., McClelland, J. L. & Hinton, G. E. (1986). Schemata and Sequential Thought Processes in PDP Models. In J. L. McClelland, D. E. Rumelhart & the PDP research Group, (Eds.), *Parallel Distributed Processing: Experiments in the Microstructure of Cognition*, Vol. 2. Cambridge: MIT Press.
- Schvaneveldt, R. W., Dearholt, D. W. & Durso, F. T. (In Press). Graph theoretic foundations of Pathfinder networks. . In *Computers and Mathematics with Applications*. .
- Shastri, L. (1988). *Semantic Nets: An evidential formalization and its connectionist realization*. Los Altos, CA: Morgan Kaufmann.
- Sumida, R. A., Dyer, M. G. & Flowers, M. (1988). Integrated Marker Passing and Connectionism for Handling Conceptual and Structural Ambiguities. In *Proceedings of the 10th Annual Conference of the Cognitive Science Society*. Montreal.
- Touretzky, D. S. & Geva, S. (1987). A distributed connectionist representation for concept structures. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*. Seattle, 155-164.
- Waltz, D. L. & Pollack, J. B. (1985). Massively Parallel Parsing: A strongly interactive model of Natural Language Interpretation. *Cognitive Science*, 9, 51-74.